arXiv:1602.05136v1 [cs.DS] 16 Feb 2016

Exploration of Faulty Hamiltonian Graphs

David Caissy * Andrzej Pelc *†

Abstract

We consider the problem of exploration of networks, some of whose edges are faulty. A mobile agent, situated at a starting node and unaware of which edges are faulty, has to explore the connected fault-free component of this node by visiting all of its nodes. The *cost* of the exploration is the number of edge traversals. For a given network and given starting node, the *overhead* of an exploration algorithm is the worst-case ratio (taken over all fault configurations) of its cost to the cost of an optimal algorithm which knows where faults are situated. An exploration algorithm, for a given network and given starting node, is called *perfectly competitive* if its overhead is the smallest among all exploration algorithms not knowing the location of faults. We design a perfectly competitive exploration algorithm for any ring, and show that, for networks modeled by hamiltonian graphs, the overhead of any DFS exploration is at most 10/9 times larger than that of a perfectly competitive algorithm. Moreover, for hamiltonian graphs of size at least 24, this overhead is less than 6% larger than that of a perfectly competitive algorithm.

Keywords: algorithm, faulty edge, exploration, mobile agent, hamiltonian graph

*Département d'informatique, Université du Québec en Outaouais, Gatineau, Québec J8X 3X7, Canada. E-mails: david.caissy@gmail.com, pelc@uqo.ca

†Research supported in part by the NSERC Discovery Grant 8136 – 2013 and by the Research Chair in Distributed Computing of the Université du Québec en Outaouais.

1 Introduction

Exploration of networks by visiting all of its nodes is one of the basic tasks performed by mobile agents in networks. In applications, a software agent may need to collect data residing at nodes of a network, or a mobile robot may need to collect samples of ground in a contaminated mine whose corridors form links of a network, with corridor crossings being nodes. In many situations, some links of the network may be faulty, preventing the mobile agent to traverse them. In a computer network a fault may be a malfunctioning link, and in a cave it may be an obstructed corridor. In this situation the agent may be unable to reach some nodes of the network, and the task is then to explore the accessible part, as efficiently as possible.

The network is modeled as a simple connected undirected graph $G = (V, E)$, called *graph* in the sequel. The agent is initially situated at a starting node v of the graph. It has a faithful labeled map of the graph with its starting position marked. Each node has a distinct label, and ports at each node of degree d are labeled by integers $1, \ldots, d$. Thus the agent knows which port at any node leads to which neighbor. However, some of the edges of the graph are faulty. The set of faulty edges $F \subseteq E$ is called a *fault configuration*. The agent does not know *a priori* the location of faults nor their number. When visiting a node for the first time, the agent discovers ports corresponding to faulty edges, if any, incident to this node. A faulty edge prevents the agent from traversing it. The fault-free subgraph G', resulting from the original graph $G = (V, E)$ after removing all faulty edges, may be disconnected. Let C be the connected component of the subgraph G', containing the starting node v. We will call C the fault-free component corresponding to the fault configuration F and to the node v. The task of the agent is to *explore* C by visiting all of its nodes. Exploration is finished when the last node of C is visited (the agent does not have to go back to its starting node). For a given graph G, starting node v, and fault configuration F, the *cost* $\mathcal{C}(A, G, v, F)$ of an exploration algorithm A is the number of edge traversals it performs when exploring the connected component C containing v and corresponding to F.

An agent that knows F, and hence knows the component C, has some algorithm exploring the component C at the minimum possible cost. Call this optimal cost $opt(G, v, F)$. Note that this optimal algorithm and its cost might be difficult to find in many cases. Now consider an exploration algorithm A for a given graph G and starting node v, that *does not know* F, as supposed in our scenario. A natural measure of performance of such an algorithm, cf. [31], is the worst-case ratio between its cost and the optimal cost $opt(G, v, F)$, where the worst case is taken over all fault configurations F. This number $\max_{F \subseteq E} \frac{\mathcal{C}(A,G,v,F)}{opt(G,v,F)}$ is called the *overhead* of A, and is denoted $\mathcal{O}_{A,G,v}$. This measure is similar to the competitive ratio of on-line algorithms. It measures the penalty incurred by the algorithm, due to some kind of ignorance. In the case of on-line algorithms, they do not know future events, which are known to an off-line algorithm (serving as a benchmark). The exploration algorithms we want to design do not know the fault configuration, which is known to an optimal algorithm (serving as a benchmark).

Given a graph G and a starting node v, an exploration algorithm (not knowing the fault configuration) is called *perfectly competitive* if it has the smallest overhead among all exploration algorithms working under this scenario. Our aim is to construct exploration algorithms with small overhead: either perfectly competitive algorithms or ones whose overhead only slightly exceeds the smallest possible overhead.

In the sequel, when the graph G and the starting node v are fixed, we will write $\mathcal{C}(A, F)$ instead of $\mathcal{C}(A, G, v, F)$, $opt(F)$ instead of $opt(G, v, F)$, and $\mathcal{O}_A$ instead of $\mathcal{O}_{A,G,v}$. If the link corresponding to port p at some node is faulty, we will say that port p at this node is *faulty*, otherwise we will say that port p is *free*.

1.1 Related work

The problem of exploration and navigation of mobile agents in an unknown environment has been extensively studied in the literature (cf. the survey [34]). Our scenario falls in this category: in our case, the unknown ingredient of the environment is the fault configuration. The explored environment has been modeled in the literature in two distinct ways: either as a geometric terrain in the plane, e.g., an unknown terrain with convex obstacles [8], or a room with polygonal [14] or rectangular [4] obstacles, or as we do, i.e., as a graph, assuming that the agent may only move along its edges. The graph model can be further specified in two different ways: either the graph is directed and strongly connected, in which case the agent can move only from tail to head of a directed edge [1, 5, 6, 16], or the graph is undirected (as we assume) and the agent can traverse edges in both directions [2, 9, 26, 32]. The efficiency measure adopted in most papers dealing with exploration of graphs is the cost of completing this task, measured by the number of edge traversals by the agent. In some papers, further restrictions on the moves of the agent are imposed. It is assumed that the robot has either a restricted tank [2, 9], and thus has to periodically return to the base for refueling, or that it is attached to the base by a rope or cable of restricted length [26].

Another direction of research concerns exploration of anonymous graphs. In this case it is impossible to explore arbitrary graphs and stop after exploration, if no marking of nodes is allowed. Hence some authors [5, 6] allow *pebbles* which the agent can drop on nodes to recognize already visited ones, and then remove them and drop them in other places. A more restrictive scenario assumes a stationary token that is fixed at the starting node of the agent [10, 33]. Exploring anonymous graphs without the possibility of marking nodes (and thus possibly without stopping) is investigated, e.g., in [19]. The authors concentrate attention not on the cost of exploration but on the minimum amount of memory sufficient to carry out this task. Exploration of anonymous graphs is also considered in [15, 24, 25].

A measure of performance of exploration algorithms, similar to the competitive ratio, and thus to our notion of overhead, has been used, e.g., in [17], to study exploration of unknown graphs and graphs for which only an unoriented map is available. In the above paper, the benchmark was the performance of an algorithm having full knowledge of the graph.

Our problem belongs to the domain of fault-tolerant graph exploration whose other aspects were studied, e.g., in [11, 12, 13, 20, 22, 23, 27, 29, 30], in the context of searching for a black hole in a network. A black hole is a process situated at an unknown node and destroying all agents visiting it. The goal is to locate a black hole using as few agents as possible, while keeping at least one agent surviving. Another aspect of fault-tolerant exploration was studied in [28]. It concerned the exploration of a line by a single agent, and faults involved directions chosen by the agent: for example, if the agent planned on going right and the fault occurred, it resulted in going left.

The paper most closely related to the present work is [31]. The authors used the same scenario

and the goal was also to find exploration algorithms with the smallest possible overhead. However, they restricted attention to the class of trees. For the simplest case of the line they designed a perfectly competitive algorithm, and they proposed another algorithm, working for arbitrary trees, whose overhead is at most 9/8 times larger than that of a perfectly competitive algorithm.

1.2 Our results

We consider hamiltonian graphs, i.e., graphs that contain a simple cycle including all nodes. This is a large class containing such important examples as complete graphs, hypercubes, rings, and even-size tori. Our goal is to design exploration algorithms for hamiltonian graphs some of whose edges are faulty, with small overhead. We have two main results.

- For any ring, and any starting node, we design a perfectly competitive exploration algorithm. The algorithm and its overhead depend only on the size of the ring. This should be contrasted with the perfectly competitive exploration algorithm for the line, designed in [31], whose behavior and overhead depend on the distance between the starting node and the closer endpoint of the line, and *do not* depend on the size of the line.
- For an arbitrary hamiltonian graph of size $n \geq 3$ and an arbitrary starting node, we show that the overhead of any DFS exploration is at most $(2n-4)/(n-1)$. We also show that this overhead is at most 10/9 times larger than that of a perfectly competitive algorithm, for any hamiltonian graph. Moreover, for hamiltonian graphs of size at least 24, this overhead turns out to be less than 6% larger than that of a perfectly competitive algorithm.

The total computation time used by our exploration algorithms is linear in the number of edges of the explored graph. Our main contribution is in the analysis of overhead of these algorithms, showing that simple and natural exploration strategies perform very well in faulty hamiltonian graphs.

2 Perfectly competitive exploration of rings

In this section we design a perfectly competitive algorithm working for an arbitrary ring. A ring is a graph $R = (V, E)$, where $V = \{v_1, ..., v_n\}$, $n \geq 3$ and $E = \{\{v_i, v_{i+1}\} : i = 1, ..., n-1\} \cup \{\{v_n, v_1\}\}$. For any $i > 1$ we call v_{i-1} the *predecessor* of v_i and for any $i < n$ we call v_{i+1} the *successor* of v_i. The predecessor of v_1 is v_n and the successor of v_n is v_1. At any node we denote by ℓ the port leading to the predecessor of v and by r the port leading to the successor of v.

We fix a ring R and a starting node v. For a fixed non-empty fault configuration F, denote by x the distance between v and the first fault encountered when always taking the port ℓ, and denote by y the distance between v and the first fault encountered when always taking the port r.

An agent knowing the configuration F, and hence knowing x and y, has the following simple exploration algorithm which is optimal: go first to the closest fault and then go back until a fault is met. Hence $opt(F)$ is $2x + y$ if $x \leq y$, and is $2y + x$ if $x > y$.

The following procedure has two parameters: a positive integer s and a *direction* $d \in \{\ell, r\}$. It causes the agent to go at most s steps always taking port d, until a fault is met. It returns the boolean value $b = false$ if a fault is met and $b = true$ otherwise.

Procedure GO-AT-DISTANCE (s, d)
    while port d at the current node is free and travelled distance is $< s$ **do**
        take port d
    if port d at the current node is free **then** $b := true$ **else** $b := false$

The next procedure causes the agent to successively take port $d \in \{\ell, r\}$ until a fault is met. It has a provision to avoid returning to the starting node v, if there are no faults.

Procedure GO-FIRM(d)
    while port d at the current node is free and fewer than n nodes have been visited **do**
        take port d

We now formulate our algorithm for ring exploration.

Algorithm `Ring`

```
if port ℓ is faulty then GO-FIRM(r)
else
    if port r is faulty then GO-FIRM(ℓ)
    else
        if n ≤ 5 then
            GO-FIRM(ℓ)
            GO-FIRM(r)
        else
            if 6 ≤ n ≤ 19 then
                GO-AT-DISTANCE(1, ℓ)
                if b then
                    GO-FIRM(r)
                    GO-FIRM(ℓ)
                else
                    GO-FIRM(r)
            else
                GO-AT-DISTANCE(2, ℓ)
                if b then
                    GO-FIRM(r)
                    GO-FIRM(ℓ)
                else
                    GO-FIRM(r)
```

Since the time of computation at each visit of a node is constant, the total computation time of Algorithm `Ring` is linear in the size of the ring. The following proposition gives the overhead of Algorithm `Ring`.

Proposition 2.1 *The overhead of Algorithm* `Ring` *is:*

- 1 *when* $n = 3$
- $\frac{2n-3}{n}$ *when* $4 \leq n \leq 5$
- $\frac{3}{2}$ *when* $6 \leq n \leq 7$
- $\frac{2n-2}{n+1}$ *when* $8 \leq n \leq 19$
- $\frac{9}{5}$ *when* $20 \leq n \leq 23$
- $\frac{2n-1}{n+2}$ *when* $n \geq 24$

Proof. First observe that if the fault configuration F is empty, then Algorithm `Ring` has cost $n-1$ if $n \leq 5$, cost n if $6 \leq n \leq 19$, and cost $n+1$ if $n \geq 20$. Hence, for the empty fault configuration F, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = 1$ if $n \leq 5$, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{n}{n-1}$ if $6 \leq n \leq 19$, and $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{n+1}{n-1}$ if $n \geq 20$. If F is non-empty then integers x and y are well defined. If $x = 0$ or $y = 0$, then $\mathcal{C}(\texttt{Ring}, F) = opt(F)$. Otherwise consider the following cases.

$n = 3$. In this case $\mathcal{C}(\texttt{Ring}, F) = 2x + y = 2y + x = opt(F)$. Hence in this case $\mathcal{O}_{\texttt{Ring}} = 1$.

$4 \leq n \leq 5$. In this case, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2x+y}{min\{2x+y,2y+x\}}$. If $x \leq y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = 1$. If $x > y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2x+y}{2y+x}$ and this fraction is maximized for $x = n-2$, $y = 1$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2n-3}{n}$. Since $1 < \frac{2n-3}{n}$ for $n \geq 4$, we get $\mathcal{O}_{\texttt{Ring}} = \frac{2n-3}{n}$ in this case.

$n = 6$. In this case, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{min\{2x+y,2y+x\}}$. If $x < y$, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{2x+y}$ and this fraction is maximized for $x = 2$, $y = 3$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{10}{7}$. If $x \geq y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{2y+x}$ and this fraction is maximized for $x = 2$, $y = 1$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{3}{2}$. The ratio $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)}$ for the empty fault configuration is $\frac{6}{5}$ in this case. Since $\frac{10}{7} < \frac{3}{2}$ and $\frac{6}{5} < \frac{3}{2}$, we get $\mathcal{O}_{\texttt{Ring}} = \frac{3}{2}$ in this case.

$n = 7$. In this case, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{min\{2x+y,2y+x\}}$. If $x \leq y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{2x+y}$ and this fraction is maximized for $x = 2$, $y = 4$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{3}{2}$. If $x > y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{2y+x}$ and this fraction is maximized for $x = 2$, $y = 1$, giving the value $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{3}{2}$. The ratio $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)}$ for the empty fault configuration is $\frac{7}{6}$ in this case. Since $\frac{7}{6} < \frac{3}{2}$, we get $\mathcal{O}_{\texttt{Ring}} = \frac{3}{2}$ in this case.

$8 \leq n \leq 19$. In this case, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{min\{2x+y,2y+x\}}$. If $x \leq y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{2x+y}$ and this fraction is maximized for $x = 2$, $y = n-3$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2n-2}{n+1}$. If $x > y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2+2y+x}{2y+x}$ and this fraction is maximized for $x = 2$, $y = 1$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{3}{2}$.

The ratio $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)}$ for the empty fault configuration is $\frac{n}{n-1}$ in this case. Since $\frac{2n-2}{n+1} > \frac{3}{2}$ and $\frac{2n-2}{n+1} > \frac{n}{n-1}$ when $n \geq 8$, we get $\mathcal{O}_{\texttt{Ring}} = \frac{2n-2}{n+1}$ in this case.

$20 \leq n \leq 23$. In this case, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{4+2y+x}{min\{2x+y,2y+x\}}$. If $x \leq y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{4+2y+x}{2x+y}$ and this fraction is maximized for $x = 3$, $y = n-4$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2n-1}{n+2}$. If $x > y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{4+2y+x}{2y+x}$ and this fraction is maximized for $x = 3$, $y = 1$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{9}{5}$. The ratio $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)}$ for the empty fault configuration is $\frac{n+1}{n-1}$ in this case. Since $\frac{2n-1}{n+2} \leq \frac{9}{5}$ and $\frac{9}{5} > \frac{n+1}{n-1}$ when $20 \leq n \leq 23$, we get $\mathcal{O}_{\texttt{Ring}} = \frac{9}{5}$ in this case.

$n \geq 24$. In this case, $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{4+2y+x}{min\{2x+y,2y+x\}}$. If $x \leq y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{4+2y+x}{2x+y}$ and this fraction is maximized for $x = 3$, $y = n-4$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{2n-1}{n+2}$. If $x > y$, then $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{4+2y+x}{2y+x}$ and this fraction is maximized for $x = 3$, $y = 1$, giving the value $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)} = \frac{9}{5}$. The ratio $\frac{\mathcal{C}(\texttt{Ring},F)}{opt(F)}$ for the empty fault configuration is $\frac{n+1}{n-1}$ in this case. Since $\frac{2n-1}{n+2} > \frac{9}{5}$ and $\frac{2n-1}{n+2} > \frac{n+1}{n-1}$ when $n \geq 24$, we get $\mathcal{O}_{\texttt{Ring}} = \frac{2n-1}{n+2}$ in this case. □

We now show that Algorithm `Ring` is perfectly competitive, i.e., has the smallest overhead among all exploration algorithms for the ring, not knowing the fault configuration. We adapt to the ring the following notions from [31] (where they were defined for the line in a similar way). We denote by $\mathcal{A}_k$ the class of exploration algorithms for the ring which do initially k returns (i.e., changes of direction) assuming that no fault is encountered before the first k returns, then GO-FIRM, and in the case when a fault is met, return and GO-FIRM. An algorithm of class $\mathcal{A}_k$ is called an i-step algorithm, for $1 \leq i \leq n-2$, if it goes i steps before the first return, unless a fault is encountered earlier. Note that $\mathcal{A}_0$ is the class of DFS algorithms (there are two such algorithms for the ring, depending on the initial direction chosen). Also, an i-step algorithm of class $\mathcal{A}_k$, for $k > 0$ and $i \geq n-1$ is in fact a DFS algorithm, i.e., it is of class $\mathcal{A}_0$. On the other hand, a 0-step algorithm of class $\mathcal{A}_k$ is an algorithm of a class $\mathcal{A}_m$, for some $m < k$, and hence it is enough to consider only i-step algorithms, for strictly positive i, in each class. Algorithm `Ring` is an algorithm of class $\mathcal{A}_0$ for $n \leq 5$, it is a 1-step algorithm of class $\mathcal{A}_1$ for $6 \leq n \leq 19$, and it is a 2-step algorithm of class $\mathcal{A}_1$ for $n \geq 20$.

The proof that Algorithm `Ring` is perfectly competitive is split into a series of lemmas.

Lemma 2.1 *Every algorithm of class $\mathcal{A}_0$ has overhead at least $\frac{2n-3}{n}$, for all $n \geq 3$.*

Proof. Consider the fault configuration F consisting of a single fault situated at distance 1 from the starting node v, in the direction different from the one at which the algorithm A of class $\mathcal{A}_0$ starts. Then $\mathcal{C}(A,F) = 2n-3$ and $opt(F) = n$, which concludes the proof. □

Lemma 2.2 *Let $n \geq 3$. Every $(n-2)$-step algorithm of class $\mathcal{A}_1$ has overhead $\frac{2n-3}{n-1}$, and every i-step algorithm of class $\mathcal{A}_1$, for $1 \leq i \leq n-3$, has overhead $max\{\frac{n+i-1}{n-1}, \frac{3i+3}{i+3}, \frac{i+2n-3}{i+n}\}$.*

Proof. First consider an $(n-2)$-step algorithm A of class $\mathcal{A}_1$, and a fault configuration F. If F is empty, then $\mathcal{C}(A,F) = 2n-3$ and $opt(F) = n-1$, giving the ratio $\frac{2n-3}{n-1}$. Otherwise, x and y are well defined. If $x = 0$ or $y = 0$, then $\mathcal{C}(A,F) = opt(F)$. Assume that $x > 0$ and $y > 0$. Without loss

of generality we may assume that in the first step the algorithm takes port ℓ. We have $x \leq n-2$, and hence $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{2x+y}{min\{2x+y,2y+x\}}$. This fraction is maximized for $x = n-2$ and $y = 1$, and then its value is $\frac{2n-3}{n}$. Since $\frac{2n-3}{n-1} > \frac{2n-3}{n}$, we get $\mathcal{O}_A = \frac{2n-3}{n-1}$.

In the rest of the proof we consider an i-step algorithm A of class $\mathcal{A}_1$, for $1 \leq i \leq n-3$, and a fault configuration F. If F is empty, then $\mathcal{C}(A,F) = n+i-1$ and $opt(F) = n-1$, giving the ratio of $\frac{n+i-1}{n-1}$. Otherwise, x and y are well defined. If $x = 0$ or $y = 0$, then $\mathcal{C}(A,F) = opt(F)$. Assume that $x > 0$ and $y > 0$. Without loss of generality we may assume that in the first step the algorithm takes port ℓ.

If $x \leq i$ then $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{2x+y}{min\{2x+y,2y+x\}}$, and this fraction is maximized for $x = i$ and $y = 1$. Hence the maximum value of $\frac{\mathcal{C}(A,F)}{opt(F)}$ is $\frac{2i+1}{i+2}$ in this case.

If $x > i$ and $x \geq y$ then $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{2i+2y+x}{2y+x}$, and this fraction is maximized for $x = i+1$ and $y = 1$. Hence the maximum value of $\frac{\mathcal{C}(A,F)}{opt(F)}$ is $\frac{3i+3}{i+3}$ in this case.

If $x > i$ and $x < y$ then $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{2i+2y+x}{2x+y}$, and this fraction is maximized for $x = i+1$ and $y = n-i-2$. Hence the maximum value of $\frac{\mathcal{C}(A,F)}{opt(F)}$ is $\frac{i+2n-3}{i+n}$ in this case.

Since $\frac{3i+3}{i+3} > \frac{2i+1}{i+2}$, il follows that $\mathcal{O}_A = \max\{\frac{n+i-1}{n-1}, \frac{3i+3}{i+3}, \frac{i+2n-3}{i+n}\}$. □

Lemma 2.2 gives the following corollary.

Corollary 2.1 *Every 1-step algorithm of class $\mathcal{A}_1$ has overhead:*

- $\frac{3}{2}$ *for* $3 \leq n \leq 7$
- $\frac{2n-2}{n+1}$ *for* $n > 7$

Every 2-step algorithm of class $\mathcal{A}_1$ has overhead:

- $\frac{5}{3}$ *for* $n = 4$
- $\frac{9}{5}$ *for* $5 \leq n \leq 23$
- $\frac{2n-1}{n+2}$ *for* $n > 23$

Notice that for $3 \leq i \leq n-3$, we have $\frac{3i+3}{i+3} \geq 2$. Moreover, $\frac{2n-3}{n-1} \geq \frac{3}{2}$ when $5 \leq n \leq 7$, and $\frac{2n-3}{n-1} \geq \frac{2n-2}{n+1}$ when $n > 7$ (i would be less than 3 for $n < 5$). Hence it follows from Lemma 2.2 and from Corollary 2.1 that the overhead of any i-step algorithm of class $\mathcal{A}_1$, for $i \geq 3$, is larger than the overhead of a 1-step algorithm of class $\mathcal{A}_1$. Hence, when looking for perfectly competitive algorithms in class $\mathcal{A}_1$ we may restrict attention to i-step algorithms for $i \leq 2$.

It was observed in [31] that if the agent, at some point of the exploration, is at node w, then moves along an already traversed edge incident to w, and immediately returns to w, then, for any fault configuration, an algorithm causing such a pair of moves has cost strictly larger than the algorithm that skips these two moves. Hence, while looking for a perfectly competitive algorithm, we may restrict attention to algorithms that do not perform such unnecessary moves. We call such

algorithms *regular*. It is easy to see that any regular exploration algorithm on a ring falls in one of the classes $\mathcal{A}_k$.

The following lemma is proved similarly as Lemma 2.5 from [31] which is its analog for the line; we include the proof for completeness. Together with the above remarks, it shows that when looking for perfectly competitive exploration algorithms on rings we may restrict attention to algorithms of class $\mathcal{A}_0$ and to 1-step and 2-step algorithms of class $\mathcal{A}_1$.

Lemma 2.3 *For every exploration algorithm of class $\mathcal{A}_k$, for $k \geq 2$, there exists an algorithm of class $\mathcal{A}_0$ or an i-step algorithm of class $\mathcal{A}_1$, for $i \leq 2$, with a smaller or equal overhead.*

Proof. Let A be any algorithm of class $\mathcal{A}_k$, for $k \geq 2$. Without loss of generality, assume that the agent starts by taking port ℓ. The behavior of A can be described as follows:

traverse z_1 edges going in one direction;

return and traverse $z_1 + z_2$ edges;

return and traverse $z_2 + z_3$ edges;

...

return and traverse $z_{k-2} + z_{k-1}$ edges;

return and traverse $z_{k-1} + z_k$ edges;

return and GO-FIRM;

return and GO-FIRM;

for $z_1 < z_3 < z_5 < ...$ and $z_2 < z_4 < z_6 < ...$, provided that the closest faults to the starting node v (if any) are at distance larger than z_{k-1} of v on the side where the agent has traversed z_{k-1} edges, and at distance larger than z_k of v on the side where the agent has traversed z_k edges. This implies that $z_{k-1} + z_k \ \leq n - 2$.

First consider a non-empty fault configuration F for which the closest fault from v in one direction is at distance $z_{k-1}+1$ from it, and the closest fault from v in the other direction is at distance z_k+1 from it (this could be the same fault). We have $\mathcal{C}(A, F) = 2(z_1 + ... + z_k) + (2z_{k-1} + 2) + (z_k + 1) = 2(z_1 + ... + z_k) + 2z_{k-1} + z_k + 3$.

If $z_k \geq z_{k-1}$, we have $opt(F) = 2z_{k-1}+z_k+3$. Hence $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{2(z_1+...+z_k)+2z_{k-1}+z_k+3}{2z_{k-1}+z_k+3}$. If additionally $2(z_1 + ... + z_k) - 2z_{k-1} - z_k \geq 3$, it follows that $\frac{\mathcal{C}(A,P)}{opt(F)} \geq 2$ in this case.

If $z_k < z_{k-1}$ then $2(z_1 + ... + z_{k-1}) \geq z_k + 3$ and $opt(A) = 2z_k + z_{k-1} + 3$. Hence $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{2(z_1+...+z_k)+2z_{k-1}+z_k+3}{2z_k+z_{k-1}+3} \geq 2$ in this case.

Since any algorithm A' of class $\mathcal{A}_0$ has an overhead $\mathcal{O}(A') < 2$, we conclude that in any of the above cases, the algorithm A' has an overhead smaller than the algorithm A.

If none of the above cases holds (i.e., if we have $z_k \geq z_{k-1}$ and $2(z_1 + ... + z_{k-2}) + z_k < 3$), then the only possibility is $k = 2$ and $z_2 \leq 2$. In this case consider a z_2-step algorithm A'' of class

$\mathcal{A}_1$. The overhead of algorithm A'' is $\mathcal{O}(A'') = \max\{\frac{n+z_2-1}{n-1}, \frac{3z_2+3}{z_2+3}, \frac{z_2+2n-3}{z_2+n}\}$ when $z_2 \leq n-3$, and $\mathcal{O}(A'') = \frac{n+z_2-1}{n-1}$ when $z_2 = n-2$, by Lemma 2.2.

We have $\frac{\mathcal{C}(A,F)}{opt(F)} = \frac{4z_1+3z_2+3}{2z_1+z_2+3} \geq \frac{3z_2+3}{z_2+3} \geq \frac{n+z_2-1}{n-1}$, because $z_2 \leq 2$.

Let F_1 be the fault configuration consisting of one fault for which $y = z_2 + 1$ and $x = n - z_2 - 2$. We have $\mathcal{C}(A, F_1) = 2z_1 + 3z_2 + 2x + 1$ and $opt(F_1) = \min\{2z_2 + x + 2, 2x + z_2 + 1\}$. Hence $\mathcal{C}(A, F_1) \geq z_2 + 2n - 3$ and $opt(F_1) \leq z_2 + n$, and thus $\frac{\mathcal{C}(A,F_1)}{opt(F_1)} \geq \frac{z_2+2n-3}{z_2+n}$.

Finally, consider the empty fault configuration $\emptyset$. We have $\mathcal{C}(A, \emptyset) = 2z_1 + z_2 + n - 1 \geq n + z_2 - 1$ and $opt(\emptyset) = n - 1$. Hence $\frac{\mathcal{C}(A,\emptyset)}{opt(\emptyset)} \geq \frac{n+z_2-1}{n-1}$.

It follows that in the case when $z_k \geq z_{k-1}$ and $2(z_1 + ... + z_{k-2}) + z_k < 3$, we have $\mathcal{O}(A'') \leq \mathcal{O}(A)$. This proves the lemma. □

Proposition 2.2 *Every exploration algorithm of the ring has overhead at least:*

- $\frac{2n-3}{n}$ *for* $4 \leq n \leq 5$
- $\frac{3}{2}$ *for* $6 \leq n \leq 7$
- $\frac{2n-2}{n+1}$ *for* $8 \leq n \leq 19$
- $\frac{9}{5}$ *for* $20 \leq n \leq 23$
- $\frac{2n-1}{n+2}$ *for* $n \geq 24$

Proof. By Lemma 2.3 we can restrict attention to algorithms of class $\mathcal{A}_0$ and i-step algorithms of class $\mathcal{A}_1$, for $i \leq 2$.

Let $4 \leq n \leq 5$. By Lemma 2.1, every algorithm of class $\mathcal{A}_0$ has overhead at least $\frac{2n-3}{n}$. By Corollary 2.1, every 1-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{3}{2}$ and every 2-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{5}{3}$. Since $\frac{2n-3}{n} \leq \frac{3}{2} < \frac{5}{3}$, the overhead of any exploration algorithm is at least $\frac{2n-3}{n}$ in this case.

Let $6 \leq n \leq 7$. By Lemma 2.1, every algorithm of class $\mathcal{A}_0$ has overhead at least $\frac{2n-3}{n}$. By Corollary 2.1, every 1-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{3}{2}$ and every 2-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{9}{5}$. Since $\frac{3}{2} \leq \frac{2n-3}{n} < \frac{9}{5}$, the overhead of any exploration algorithm is at least $\frac{3}{2}$ in this case.

Let $8 \leq n \leq 19$. By Lemma 2.1, every algorithm of class $\mathcal{A}_0$ has overhead at least $\frac{2n-3}{n}$. By Corollary 2.1, every 1-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{2n-2}{n+1}$ and every 2-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{9}{5}$. Since $\frac{2n-2}{n+1} < \frac{2n-3}{n} \leq \frac{9}{5}$ for $8 \leq n \leq 15$ and $\frac{2n-2}{n+1} \leq \frac{9}{5} < \frac{2n-3}{n}$ for $16 \leq n \leq 19$, the overhead of any exploration algorithm is at least $\frac{2n-2}{n+1}$ in this case.

Let $20 \leq n \leq 23$. By Lemma 2.1, every algorithm of class $\mathcal{A}_0$ has overhead at least $\frac{2n-3}{n}$. By Corollary 2.1, every 1-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{2n-2}{n+1}$ and every 2-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{9}{5}$. Since $\frac{9}{5} < \frac{2n-2}{n+1} < \frac{2n-3}{n}$, the overhead of any exploration algorithm is at least $\frac{9}{5}$ in this case.

Let $n \geq 24$. By Lemma 2.1, every algorithm of class $\mathcal{A}_0$ has overhead at least $\frac{2n-3}{n}$. By Corollary 2.1, every 1-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{2n-2}{n+1}$ and every 2-step algorithm of class $\mathcal{A}_1$ has overhead at least $\frac{2n-1}{n+2}$. Since $\frac{2n-1}{n+2} < \frac{2n-2}{n+1} < \frac{2n-3}{n}$, the overhead of any exploration algorithm is at least $\frac{2n-1}{n+2}$ in this case. □

Propositions 2.1 and 2.2 imply:

Theorem 2.1 *Algorithm* `Ring` *is perfectly competitive for an arbitrary ring.*

Remark. It is interesting to compare our scenario of exploration of an n-node ring with possibly faulty links to the scenario from [17] of exploration of a line of known size n (without faults) by an agent starting in an unknown node of the line. (This scenario was called *exploration with an unanchored map* in [17]). Notice that the latter scenario is equivalent to our ring scenario with exactly one fault. In [17], an algorithm with optimal overhead was given and proved to have overhead $\sqrt{3}$. Since there were no faults in [17], the overhead involved a maximum ratio over all positions of the starting node. This algorithm, if applied in our case of rings with possibly many faulty links, would have a huge overhead: arbitrarily close to 3, for large n, and thus much worse than DFS. This is due to the fact that, in the scenario from [17], the adversary can only place the starting node arbitrarily, but has no power of changing the size of the line, as this is known to the agent. By contrast, in a ring with arbitrary fault configurations, the adversary can place arbitrarily the two faults that define the size of the connected component containing the starting node. This is what makes the optimal algorithm from [17] so bad in our present scenario.

3 DFS exploration of hamiltonian graphs

Since the agent has a faithful map of the graph, for most graphs there are many DFS (Depth-First-Search) exploration algorithms of the fault-free component of the graph, depending on the order in which free ports at each visited node are explored. Thus it is appropriate to speak of algorithms of the class $\mathcal{DFS}$, each of whose members is specified by giving a permutation of ports at each node that determines this order. Hence we adopt the following definition.

Fix an arbitrary graph $G = (V, E)$ and any fault configuration F in G. Let v be the starting node of the agent. For any node w of G, let σ_w be a permutation of ports at w. Let $\alpha = (\sigma_w : w \in V)$. $DFS(\alpha)$ is the algorithm specified by calling the following recursive procedure at node v.

Procedure `Explore`(w)
Let $(i_1, \ldots, i_k)$ be the sequence of free ports at w listed in order given by σ_w.
Let $(v_1, \ldots, v_k)$ be the list of neighbors of w corresponding to these ports.
For any $r = 1, \ldots, k$, let j_r be the port at v_r corresponding to the edge $\{v_r, w\}$.
mark(w)
for $r = 1$ **to** k **do**
    if v_r is unmarked **then**
        take port i_r
        `Explore`(v_r)
        if there are still unmarked nodes **then** take port j_r

The class $\mathcal{DFS}$ is defined as the class of algorithms $DFS(\alpha)$, for all sequences of permutations $\alpha = (\sigma_w : w \in V)$. Note that for any algorithm of the class $\mathcal{DFS}$, the route of the agent is a tour of some spanning tree of the fault-free component C of the graph, containing the starting node and corresponding to the fault configuration F, in which the agent stops as soon as the last node of the component is visited. Hence the edge of the tree corresponding to the leaf that is visited last is traversed only once. Consequently, for any α and any F, we have $\mathcal{C}(DFS(\alpha), F) \leq 2m - 3$, where m is the size of the component C. The total computation time of any algorithm of the class $\mathcal{DFS}$ is linear in the number of all ports, i.e., also linear in the number of edges of the explored graph.

Our aim is to show that if the graph G is hamiltonian then any algorithm of the class $\mathcal{DFS}$ has overhead not much larger than that of a perfectly competitive algorithm. Note that a hamiltonian graph must have at least three nodes, and hence we assume in the sequel that the size n of the graph is at least 3. We first prove the following lemma.

Lemma 3.1 *Let G be any hamiltonian graph of size $n \geq 3$. For any algorithm A of the class $\mathcal{DFS}$ and any starting node, $\mathcal{O}_A \leq \frac{2n-4}{n-1}$.*

Proof. Consider any fault configuration F and any starting node v. Let C be the fault-free component, corresponding to the configuration F and to the node v. Let m be the size of C. Consider two cases.

Case 1. There exists a hamiltonian path in C with v as an endpoint.

Let $(w_1, \ldots, w_m)$, where $w_1 = v$, be a hamiltonian path in C. C contains all edges of this path and possibly some additional edges $\{w_i, w_j\}$. For $m = 2$, we have $\mathcal{C}(A, F) = opt(F) = 1$. Assume $m \geq 3$; we have $opt(F) = m - 1$. Let T be the spanning tree of C, corresponding to the route of the agent executing algorithm A. T has $m - 1$ edges. As previously remarked, the last-traversed edge of T is traversed only once. We show that the first-traversed edge is also traversed only once. Let $\{w_1, w_j\}$ be this edge. Suppose that this edge is traversed by the agent a second time. This traversal must be from w_j to w_1. This implies that at the time t of the visit of w_j immediately preceding this traversal, there are still non-visited nodes. Let w_s be the node non-visited at time t, with the largest index $s < j$ or with the smallest index $s > j$. If $s < j$ then w_{s+1} was visited before time t and if $s > j$ then w_{s-1} was visited before time t. In the first case the port at w_{s+1} corresponding to the edge $\{w_{s+1}, w_s\}$ is free, and in the second case the port at w_{s-1} corresponding to the edge $\{w_{s-1}, w_s\}$ is free. This contradicts the fact that, according to algorithm A, the agent should visit w_s before time t.

Hence both the first-traversed edge and the last-traversed edge are traversed only once. This implies that the cost of algorithm A is at most $2m - 4$, for $m \geq 3$, in this case.

Case 2. There is no hamiltonian path in C with v as an endpoint.

In this case $opt(F) \geq m$ and the cost of algorithm A is at most $2m - 3$ because the last-traversed edge is traversed only once.

Since $\frac{2m-4}{m-1} \geq \frac{2m-3}{m}$ for $m \geq 3$, it follows that in both cases $\frac{\mathcal{C}(A,F)}{opt(F)} \leq \frac{2m-4}{m-1}$, whenever the component C is of size m. Since $\frac{2m-4}{m-1}$ is an increasing function of m, this value is largest for $m = n$, which concludes the proof. □

The following is the main result of this section.

Theorem 3.1 *For any hamiltonian graph $G = (V, E)$, any starting node v, and any algorithm A of the class $\mathcal{DFS}$, the ratio between the overhead of A and the overhead of a perfectly competitive algorithm on G starting at v is at most 10/9. Moreover, for $n \geq 24$, this ratio is less than 1.06.*

Proof. Let B be a perfectly competitive algorithm on G, starting at v. Let R be any hamiltonian cycle in G. By definition, R contains the starting node v. Let F^* be the set of all edges in G except those in R. Let $x = \max_{F^* \subseteq F \subseteq E} \frac{\mathcal{C}(B,G,v,F)}{opt(G,v,F)}$. By definition, $\mathcal{O}_{B,G,v} \geq x$. On the other hand, the execution of any algorithm for a fault configuration F containing F^* can be considered as an execution of an algorithm on the ring R for the fault configuration $F \setminus F^*$. Since the family of fault configurations in G containing F^* is equal to the family of fault configurations $F' \cup F^*$, where F' is a fault configuration in R, we get $x \geq \mathcal{O}_{\texttt{Ring},R,v}$, for otherwise Algorithm `Ring` would not be perfectly competitive on R, contradicting Theorem 2.1. It follows that $\mathcal{O}_{B,G,v} \geq \mathcal{O}_{\texttt{Ring},R,v}$. By Proposition 2.2, $\mathcal{O}_{B,G,v}$ is at least:

- $\frac{2n-3}{n}$ for $4 \leq n \leq 5$
- $\frac{3}{2}$ for $6 \leq n \leq 7$
- $\frac{2n-2}{n+1}$ for $8 \leq n \leq 19$
- $\frac{9}{5}$ for $20 \leq n \leq 23$
- $\frac{2n-1}{n+2}$ for $n \geq 24$

By Lemma 3.1, the ratio $\frac{\mathcal{O}_{A,G,v}}{\mathcal{O}_{B,G,v}}$ is at most:

- 1 when $n = 3$
- $\frac{\frac{2n-4}{n-1}}{\frac{2n-3}{n}} \leq \frac{15}{14} \approx 1.0714$ when $4 \leq n \leq 5$
- $\frac{\frac{2n-4}{n-1}}{\frac{3}{2}} \leq \frac{10}{9} \approx 1.1111$ when $6 \leq n \leq 7$
- $\frac{\frac{2n-4}{n-1}}{\frac{2n-2}{n+1}} \leq \frac{54}{49} \approx 1.1020$ when $8 \leq n \leq 19$
- $\frac{\frac{2n-4}{n-1}}{\frac{9}{5}} \leq \frac{35}{33} \approx 1.0606$ when $20 \leq n \leq 23$
- $\frac{\frac{2n-4}{n-1}}{\frac{2n-1}{n+2}} \leq \frac{1144}{1081} \approx 1.0583$ when $n \geq 24$

Since $1 < \frac{1144}{1081} < \frac{35}{33} < \frac{15}{14} < \frac{54}{49} < \frac{10}{9}$, the ratio between the overhead of A and the overhead of a perfectly competitive algorithm on G starting at v is at most $10/9$. Moreover, for $n \geq 24$, this ratio is less than 1.06. □

It is natural to ask if the bound $\frac{2n-4}{n-1}$ on the overhead of any algorithm of the class $\mathcal{DFS}$, obtained in Lemma 3.1, is tight. This is not the case for all hamiltonian graphs, as witnessed by the example of the ring. By Lemma 2.1, the overhead of any algorithm of the class $\mathcal{DFS}$ on the ring is $\frac{2n-3}{n} < \frac{2n-4}{n-1}$. Hence a natural question is whether this bound is tight for some hamiltonian graphs. The next proposition shows that the answer to this question is positive.

Proposition 3.1 *The overhead of any algorithm A of the class $\mathcal{DFS}$ is exactly $\frac{2n-4}{n-1}$ for the complete graph of size $n \geq 3$.*

Proof. Let A be any algorithm of the class $\mathcal{DFS}$ starting at a node v of the complete graph K_n. In view of Lemma 3.1, it is enough to show a fault configuration F, for which $\frac{\mathcal{C}(A,K_n,v,F)}{opt(K_n,v,F)} = \frac{2n-4}{n-1}$. This configuration depends on the set of permutations $\alpha = (\sigma_w : w \in V)$ for which $A = DFS(\alpha)$. Let $\sigma_v = (i_1, \ldots, i_{n-1})$ and let w be the node corresponding to port i_1. Let $\sigma_w = (j_1, \ldots, j_{n-1})$ and let j_k be the port number at w corresponding to the edge $\{w, v\}$. Let $a = \min\{1, \ldots, n-1\} \setminus \{k\}$ and let $b = \min\{1, \ldots, n-1\} \setminus \{k, a\}$. Let u' be the node corresponding to port j_a at node w and let u'' be the node corresponding to port j_b at node w. We define $v_1 = v$, $v_{n-2} = u'$, $v_{n-1} = w$, $v_n = u''$, and all nodes other than v, u', w, u'' are called arbitrarily $v_2, \ldots, v_{n-3}$. Let P be the set of edges $\{\{v_1, v_{n-1}\}, \{v_1, v_2\}, \{v_2, v_3\}, \ldots, \{v_{n-1}, v_n\}\}$ (see Fig. 1).

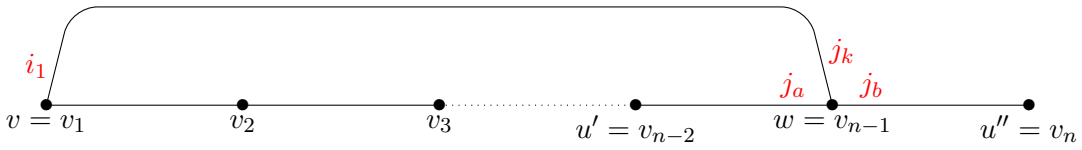

Figure 1: The set of edges P

Consider the fault configuration F consisting of all edges of K_n except those in P. For this fault configuration the optimal cost is $n-1$, corresponding to traveling the path $v_1, v_2, \ldots, v_n$. The agent executing algorithm A first takes port i_1 going to node w, then takes port j_a which is the first free port at w according to permutation σ_w, apart from the port j_k by which the agent has just come to w. At this point the agent is at $u' = v_{n-2}$. Next the agent travels the path $(v_{n-2}, v_{n-3}, \ldots, v_2)$, backtracks by the reverse path $(v_2, v_3, \ldots, v_{n-2})$ to $u' = v_{n-2}$, backtracks to $w = v_{n-1}$, and finally goes to $u'' = v_n$. The total cost is $2n-4$. Hence $\frac{\mathcal{C}(A,K_n,v,F)}{opt(K_n,v,F)} = \frac{2n-4}{n-1}$. □

4 Conclusion

We designed a perfectly competitive exploration algorithm for rings, and we proved that, for all hamiltonian graphs, the overhead of *every* algorithm of the class $\mathcal{DFS}$ exceeds that of a perfectly competitive algorithm for the given graph by less than a factor of $10/9$. The latter result does not hold for arbitrary graphs. This is best witnessed by the case of a line, for which the starting node is at distance 1 from one of the endpoints. In this case the perfectly competitive algorithm has

overhead 1 (as shown in [31]), and the "wrong" DFS (i.e., the one that goes first to the farther endpoint) has overhead $\frac{2n-1}{n+1}$, for lines of length n, i.e., it is arbitrarily close to 2 for long lines. This still leaves the hope that the "best" member of the class $\mathcal{DFS}$, i.e., the DFS algorithm with the smallest overhead, for any given graph, may have overhead similar to that of a perfectly competitive algorithm for this graph. This might be true, but finding this best DFS (a trivial task in the case of the line) may be very challenging for arbitrary graphs. We know neither if the "best" DFS for an arbitrary graph has an overhead similar to that of a a perfectly competitive algorithm for this graph, nor, even if this is the case, how to find this particular DFS, or any other of similar overhead.

Going beyond the class $\mathcal{DFS}$, the most important question seems to be, whether there exists a perfectly competitive algorithm for arbitrary graphs, whose total computation time is polynomial in the size of the graph. Unless such an algorithm is constructed, an important challenge would be finding, for an arbitrary graph, some exploration algorithm whose overhead exceeds that of a perfectly competitive algorithm for the given graph by a small factor. In view of the present paper and of [31] the situation is the following. For the class of hamiltonian graphs and for the class of trees, such algorithms exist, and the ratio is at most 10/9 in the first case and at most 9/8 in the second case. For these two classes of graphs, the algorithms performing so well are very simple, and the total computation time used during exploration is linear in the number of edges of the explored graph. (Since *any* algorithm of the class $\mathcal{DFS}$ on hamiltonian graphs is good for our purpose, we can take, e.g., the one induced by the increasing permutation of ports at each node.) Is it possible to design, for an arbitrary graph, an exploration algorithm whose overhead exceeds that of a perfectly competitive algorithm for this graph by a factor of, say, at most 1.15, and for which the total computation time is polynomial in the size of the graph?

Other interesting variants of the problem of exploration in the presence of faulty edges are when the agent knows the exact value of the number of faults, or some upper bound on this number.

References

[1] S. Albers and M. R. Henzinger, Exploring unknown environments, SIAM Journal on Computing 29 (2000), 1164-1188.

[2] B. Awerbuch, M. Betke, R. Rivest and M. Singh, Piecemeal graph learning by a mobile robot, Information and Computation 152 (1999), 155-172.

[3] B. Balamohan, P. Flocchini, A. Miri and N. Santoro, Time optimal algorithms for black hole search in rings, Discrete Math., Alg. and Appl. 3 (2011), 457-472.

[4] E. Bar-Eli, P. Berman, A. Fiat and R. Yan, On-line navigation in a room, Journal of Algorithms 17 (1994), 319-341.

[5] M.A. Bender, A. Fernandez, D. Ron, A. Sahai and S. Vadhan, The power of a pebble: Exploring and mapping directed graphs, Information and Computation 176 (2002), 1-21.

[6] M.A. Bender and D. Slonim, The power of team exploration: Two robots can learn unlabeled directed graphs, Proc. 35th Ann. Symp. on Foundations of Computer Science (FOCS 1994), 75-85.

[7] P. Berman, A. Blum, A. Fiat, H. Karloff, A. Rosen and M. Saks, Randomized robot navigation algorithms, Proc. 7th ACM-SIAM Symp. on Discrete Algorithms (SODA 1996), 74-84.

[8] A. Blum, P. Raghavan and B. Schieber, Navigating in unfamiliar geometric terrain, SIAM Journal on Computing 26 (1997), 110-137.

[9] M. Betke, R. Rivest and M. Singh, Piecemeal learning of an unknown environment, Machine Learning 18 (1995), 231-254.

[10] J. Chalopin, S. Das and A. Kosowski, Constructing a map of an anonymous graph: Applications of universal sequences, Proc. 14th International Conference on Principles of Distributed Systems (OPODIS 2010), 119-134.

[11] C. Cooper, R. Klasing and T. Radzik, Searching for black-hole faults in a network using multiple agents, Proc. 10th International Conference on Principles of Distributed Systems (OPODIS 2006), 320–332.

[12] C. Cooper, R. Klasing and T. Radzik, Locating and repairing faults in a network with mobile agents, Theoretical Computer Science, 411 (2010), 1638–1647.

[13] J. Czyzowicz, D. Kowalski, E. Markou and A. Pelc, Searching for a black hole in synchronous tree networks, Combinatorics, Probability & Computing 16 (2007), 595-619.

[14] X. Deng, T. Kameda and C. H. Papadimitriou, How to learn an unknown environment I: the rectilinear case, Journal of the ACM 45 (1998), 215-245.

[15] X. Deng and A. Mirzaian, Competitive robot mapping with homogeneous markers, IEEE Transactions on Robotics and Automation 12 (1996), 532-542.

[16] X. Deng and C. H. Papadimitriou, Exploring an unknown graph, Journal of Graph Theory 32 (1999), 265-297.

[17] A. Dessmark and A. Pelc, Optimal graph exploration without good maps, Theoretical Computer Science 326 (2004), 343-362.

[18] Y. Dieudonné and A. Pelc, Deterministic network exploration by a single agent with Byzantine tokens, Information Processing Letters 112 (2012), 467-470.

[19] K. Diks, P. Fraigniaud, E. Kranakis and A. Pelc, Tree exploration with little memory, Journal of Algorithms 51 (2004), 38-63.

[20] S. Dobrev, P. Flocchini, R. Kralovic, P. Ruzicka, G. Prencipe and N. Santoro, Black hole search in common interconnection networks, Networks 47 (2006), 61-71.

[21] S. Dobrev, P. Flocchini, R. Kralovic and N. Santoro, Exploring an unknown dangerous graph using tokens, Theoretical Computer Science 472 (2013), 28-45.

[22] S. Dobrev, P. Flocchini, G. Prencipe, and N. Santoro, Mobiles search for a black hole in an anonymous ring, Algorithmica 48 (2007), 67-90.

[23] S. Dobrev, P. Flocchini, G. Prencipe and N. Santoro, Searching for a black hole in arbitrary networks: Optimal mobile agents protocols, Distributed Computing 19 (2006), 1-18.

[24] G. Dudek, M. Jenkin, E. Milios and D. Wilkes, Robotic exploration as graph construction, IEEE Transactions on Robotics and Automation 7 (1991), 859-865.

[25] V. Dujmović and Sue Whitesides, On validating planar worlds, Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2001), 791-792.

[26] C.A. Duncan, S.G. Kobourov and V.S.A. Kumar, Optimal constrained graph exploration, ACM Transactions on Algorithms 2 (2006), 380-402.

[27] P. Flocchini, M. Kellett, P. C. Mason and N. Santoro, Searching for black holes in subways, Theory of Computing Systems 50 (2012), 158-184.

[28] P. Flocchini, A. Pelc and N. Santoro, Fault-tolerant sequential scan, Theory of Computing Systems 45 (2009), 1-26.

[29] R. Klasing, E. Markou, T. Radzik and F. Sarracco, Hardness and approximation results for black hole search in arbitrary graphs, Theoretical Computer Science 384 (2007), 201–221.

[30] R. Klasing, E. Markou, T. Radzik and F. Sarracco, Approximation bounds for Black Hole Search problems, Networks 52 (2008), 216–226.

[31] E. Markou and A. Pelc, Efficient exploration of faulty trees, Theory of Computing Systems 40 (2007), 225-247.

[32] P. Panaite and A. Pelc, Exploring unknown undirected graphs, Journal of Algorithms 33 (1999), 281-295.

[33] A. Pelc and A. Tiane, Efficient grid exploration with a stationary token, International Journal of Foundations of Computer Science 25 (2014), 247-262.

[34] N. S. V. Rao, S. Kareti, W. Shi and S.S. Iyengar, Robot navigation in unknown terrains: Introductory survey of non-heuristic algorithms, Tech. Report ORNL/TM-12410, Oak Ridge National Laboratory, July 1993.